The Track Repulsion Effect in Automatic Tracking

Stefano Coraluppi, Craig Carthel
Applied Research Department
NATO Undersea Research Centre
Viale S. Bartolomeo 400, 19126 La Spezia, Italy
coraluppi@nurc.nato.int, carthel@nurc.nato.int

Peter Willett
University of Connecticut – ECE Department
Storrs, CT 06269, USA
willett@engr.uconn.edu

Maxence Dingboe, Owen O'Neill
French Naval Academy
29240 Brest Armees, France
maxence.dingboe@hotmail.fr, owen.oneill@hotmail.fr

Tod Luginbuhl
Naval Undersea Warfare Center
Newport, RI 02841, USA
t.e.luginbuhl@ieee.org

Abstract –*The track coalescence effect degrades the performance of probabilistic data association trackers in dense target scenarios. Recently, it has been observed that an opposite effect exists with trackers that utilize hard data association, which we denote as the track repulsion effect. In this paper, we examine this effect in the context of a crossing target scenario, and explore the effectiveness of a track-oriented multi-hypothesis tracker in combating this effect, with both single-stage and multi-stage processing configurations.*

Keywords: Target tracking, multi-hypothesis tracking, distributed processing, track coalescence, track repulsion effect.

1 Introduction

Historically, military surveillance research has focused heavily on sensor technology. Downstream sensor fusion and target tracking technology has received less attention, and is an area where considerable performance gains remain to be achieved.

A broad overview of approaches to data fusion is provided in [1]. Some approaches are appropriate for expeditionary operations that do not require real-time surveillance; as an example, area clearance prior to passage of a high-value unit requires surveillance results at the end of the data acquisition period. This allows for powerful batch-processing methods to be brought to bear on the problem [2]. On the other hand, scan-based methods must be utilized for real-time surveillance. Optimal data fusion remains a *holy grail* of sorts, in that all proposed fusion paradigms are known to invoke a number of simplifying algorithmic assumptions.

The most powerful current approach to data fusion is multi-hypothesis tracking, which was first introduced in the late 1970s [3] and made feasible in the mid-1980s with the track-oriented approach [4]. A number of enhancements to the basic approach have appeared over the years [1].

The NURC *distributed multi-hypothesis tracker* (DMHT) is a high-performance, computationally efficient, and modular algorithm that was developed for undersea surveillance with a network of active sonar systems [5] and is being extended in support of the NURC *Maritime Surveillance System* [6]. The unifying theme for much of our research in data fusion has been the following: high-performance tracking requires an effective choice of multi-stage data fusion architecture. Indeed, in specific settings, multi-stage fusion is shown to outperform single-stage, centralized, *track-while-fuse* processing. The reader is directed to [7] and references therein for further details on a number of multi-stage architectures and applications: *track-before-fuse* (ground and undersea domains), *fuse-before-track* (large sensor fields), *track-extract-track* (high-clutter maritime domain), and *track-break-track* (difficult multi-target scenarios). The specific architecture that will be relevant in this paper is *track-break-track*. In sections 2-4, we provide a brief overview to the DMHT,

		Form Approved OMB No. 0704-0188

Report Documentation Page

1. REPORT DATE **JUL 2009**	2. REPORT TYPE	3. DATES COVERED **06-07-2009 to 09-07-2009**

4. TITLE AND SUBTITLE **The Track Repulsion Effect in Automatic Tracking**		5a. CONTRACT NUMBER
		5b. GRANT NUMBER
		5c. PROGRAM ELEMENT NUMBER
6. AUTHOR(S)		5d. PROJECT NUMBER
		5e. TASK NUMBER
		5f. WORK UNIT NUMBER

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) **Applied Research Department,NATO Undersea Research Center,Viale S. Bartolomeo 400, 19126 La Spezia, Italy, ,**	8. PERFORMING ORGANIZATION REPORT NUMBER
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)	10. SPONSOR/MONITOR'S ACRONYM(S)
	11. SPONSOR/MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION/AVAILABILITY STATEMENT
Approved for public release; distribution unlimited

13. SUPPLEMENTARY NOTES
See also ADM002299. Presented at the International Conference on Information Fusion (12th) (Fusion 2009). Held in Seattle, Washington, on 6-9 July 2009. U.S. Government or Federal Rights License.

14. ABSTRACT
The track coalescence effect degrades the performance of probabilistic data association trackers in dense target scenarios. Recently, it has been observed that an opposite effect exists with trackers that utilize hard data association, which we denote as the track repulsion effect. In this paper, we examine this effect in the context of a crossing target scenario, and explore the effectiveness of a track-oriented multi-hypothesis tracker in combating this effect, with both single-stage and multistage processing configurations.

15. SUBJECT TERMS

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT **Public Release**	18. NUMBER OF PAGES **6**	19a. NAME OF RESPONSIBLE PERSON
a. REPORT **unclassified**	b. ABSTRACT **unclassified**	c. THIS PAGE **unclassified**			

Standard Form 298 (Rev. 8-98)
Prescribed by ANSI Std Z39-18

the *track-break-track* architecture, and tracker performance evaluation.

The *track coalescence effect* degrades the performance of probabilistic data association trackers in dense target scenarios [8]; attempts at sub-optimal Bayesian processing to combat this effect have been reported [9]. Recently, it has been observed that an opposite effect exists with trackers that utilize hard data association, which we denote as the *track repulsion effect* [10]. In section 5, we examine this effect in the context of a crossing target scenario, and explore the effectiveness of a track-oriented multi-hypothesis tracker in combating this effect with both single-stage and multi-stage processing. Further analysis is provided in section 6, and section 7 gives conclusions and recommendations for further work.

2   A Multi-Stage Multi-Hypothesis Tracking Approach

The DMHT is a computationally-efficient, high-performance, flexible multi-hypothesis tracking approach that enables multi-stage fusion processing. Our exploitation of this tool for challenging surveillance problems is ongoing, and includes non-military applications [11].

We now illustrate the basic track-oriented *multi-hypothesis tracking* (MHT) approach with a simple example, shown in figure 1.
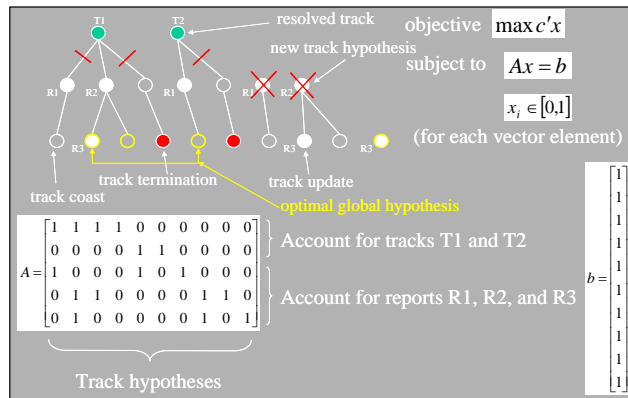

Figure 1. A simple MHT example.

The example assumes that two tracks, T1 and T2, have already been *resolved*. That is, prior data association decisions have led to a single global hypothesis that includes two tracks. Next, assume that a scan of data is received with two measurements, R1 and R2. Assume further that both R1 and R2 can feasibly be associated with T1, while only R1 can feasibly be associated with T2. This leads to a number of *local* (or *track*) hypotheses. Note that this set of hypotheses includes track continuation in the absence of a measurement (often denoted a track *coast*), as well as new-track hypotheses. A second scan of data includes a single measurement R3. We assume that R3 provides feasible updates to track

hypotheses that include R2, as well as spawning a new-track hypothesis. Note that we assume that tracks are terminated after two coasts, indicated by the red icons in figure 1.

While the example includes a number of track hypotheses, it is important to note that each *global* hypothesis provide a compete set of data-association decisions that account for all resolved tracks and all sensor measurements. The number of global hypotheses is large, even for this simple example; the power of the track-oriented approach is that we do not require an explicit enumeration of global hypotheses.

Each track hypothesis has an associated log-likelihood score that reflects track initiation and termination penalties as well as nonlinear filtering scoring; in the case of linear Gaussian systems, this scoring is based on the filter *innovations* [12]. The vector c includes the track-hypothesis scores. We are interested in the optimal global hypothesis, which amounts to identifying a vector x such that the global log-likelihood is maximized: the *maximum likelihood* solution. Having identified this solution through a two-stage relaxation approach based on linear programming or Lagrangian relaxation [13] (solution is noted in yellow in figure 1), many *conflicting* local hypotheses are removed. In particular, those track hypotheses that differ in the first scan past the resolved hypothesis layer are removed, while those that differ in the more recent past are maintained.

Having pruned the set of track hypothesis trees (with 5 surviving track hypotheses), we are ready for a new scan of data. In the example, the resolved layer always lags the current time by one scan: thus we have a multi-hypothesis example with hypothesis-tree depth (*n-scan*) of one.

Multi-stage fusion with the DMHT has two defining characteristics that differ from many legacy systems that exist today [1]. The first is that each tracker module retains measurement-level information at the output. That each, each module performs the following: it removes large numbers of measurement data, and associates the remaining measurements to form tracks over time. If the tracker is working well and the data is of reasonable quality, false measurements will be largely removed, and target-originated measurements will be largely maintained, and associated into tracks that persist over time with limited fragmentation. Since measurement data is available at the tracker output, optimal track fusion and state estimation is achievable in downstage tracker modules; the cost to achieve this performance benefit is a slightly larger bandwidth requirement between processing stages. The second defining characteristic of the DMHT is that track fusion is achieved in real time, with a scan-based approach. Traditionally, track fusion is performed in a post-processing batch mode that is not readily amenable to real-time surveillance application [1].

The theoretical optimality of unified, batch and centralized approaches to fusion and tracking (*track-while-fuse*) is at odds with a number of practical

considerations. First, in many surveillance settings optimal processing algorithms are either not known, or are computationally infeasible. Second, detection-level data may not be available from some propriety or legacy sensor systems; thus, in general it may be required to process a mix of track-level and measurement-level data. Finally, as we will see in subsequent sections, improved performance can be achieved with multi-stage processing that involves simpler and less computationally intensive algorithms than with near-optimal centralized processing.

The DMHT provides an ideal tool to explore the superior performance that can be achieved with distributed, multi-stage algorithms. Many of the findings are seemingly at odds with fundamental results in the nonlinear filtering and distributed detection literature, and are based on the fundamental sub-optimality of all current approaches to target tracking that must contend with data-association uncertainty.

3 The Track-Break-Track Architecture

We introduce a novel approach based on a *track-break-track* architecture that leverages the modularity in the DMHT. Specifically, we perform a first stage of tracking with *n-scan*=0; this often results in track swapping or other undesirable tracking effects. The value of the first tracking stage is that it removes significant numbers of extraneous contacts. Next, we break all contact associations, and provide the resulting cleaner set of contact data (with FAR close to zero) to a second tracking stage, now with *n-scan*>0. An illustration of the *track-break-track* architecture is in figure 2.
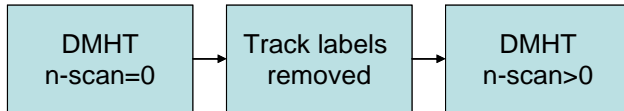

Figure 2. The *track-break-track* architecture.

We will see that the DMHT in a *track-break-track* configuration is effective at combating the track repulsion effect. In particular, computation times are significantly reduced, allowing for large values of the hypothesis tree depth (*n-scan*). Further, tracking results outperform those based on a single-stage approach.

Another application of the *track-break-track* configuration in the context of simulated active sonar datasets is in [7].

4 Tracker Performance Evaluation

Tracker performance evaluation [14] requires that, subject to a global (average) localization threshold, tracks be partitioned as target-originated or not. The former set then provides the basis for the computation of *track hold* or *track PD* (track duration as fraction of target duration), *fragmentation rate* (ratio of number of true tracks to targets per unit time), and *localization error* (average

discrepancy in meters between track positional estimate and target location). The latter set of tracks provides the basis for the *false track rate* or *FTR* (average number of false tracks per unit time).

These metrics are insufficient to capture track-swap phenomena illustrated in figure 3. Thus, we introduce an iterative scheme that breaks tracks when these are not consistently mapped to the same target. When a sufficient number of consecutive track updates is associated with a different target (or with none) relative to the global mapping of most frequent mappings, a track break is introduced. The methodology relies on a distance mapping threshold and a maximum-anomaly threshold. Correspondingly, an additional metrics is the *broken track rate* or BTR, i.e. the number of breaks per unit time.

The BTR metric is crucial, as it reflects the extent of tracker output manipulation prior to performance evaluation. Further, a BTR>0 is an indication that a track swap has occurred.
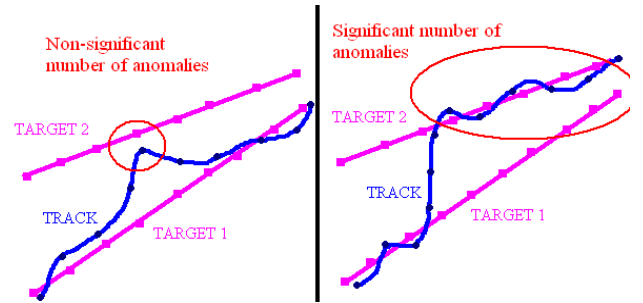

Figure 3. Track break (right) in response to a significant track-truth mapping anomaly.

5 The Track Repulsion Effect

As previously observed in the literature [10], automatic trackers exhibit a track repulsion effect whereby neighbouring targets lead to tracks that are displaced at greater distances than the targets themselves. For targets that approach slowly, this displacement leads to track swapping. We study this phenomenon with DMHT. In particular, we use the following experimental setup:

- Ground truth: two constant velocity targets, x velocity = 500m / 179sec, y velocity = +/- x velocity · tan(target angle) / 2;
- Contact data: 180 scans of data; 1sec scan repetition time; PD=1, FAR=0 (ideal case), PD=0.9, FAR=7 (non-ideal case); positional measurements with std. dev. error of 1m/s in both x and y;
- Automatic tracker (DMHT) settings: process noise $q=0.01m^2/s^3$;
- Performance assessment methodology: starting at 30deg, decrease target angle until track swap is observed: this defines the *critical angle*;
- Results are averaged over 50 Monte Carlo realizations.

Illustrations of the scenario and of the track swapping phenomenon are in figures 4-5.
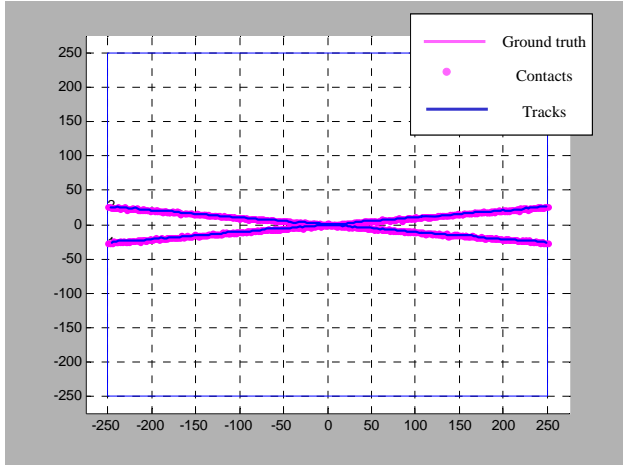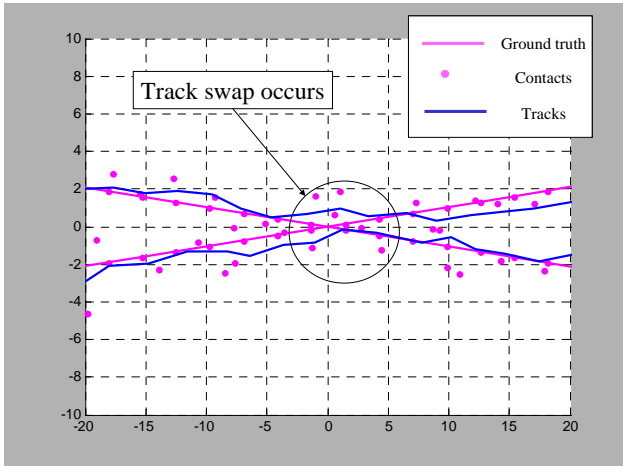

Figure 4. One realization (wide view).

Figure 5. One realization (close-up view).

We first study the ideal case of detection data with PD=1 and FAR=0. The resulting critical angle as a function of the hypothesis depth in multi-hypothesis processing, or *n-scan*, is given in figure 6. We see that the impact of the phenomenon decreases as we increase the effectiveness of the tracking algorithm, though this comes at increasing computational expense. Further, we note that the performance benefits saturate beyond *n-scan*=3.

For the general case (PD<1, FAR>0), a more effective approach to combating the track swap phenomenon is required. We introduce a novel approach based on a *track-break-track* architecture that leverages the modularity in the DMHT. Specifically, we perform a first stage of tracking with *n-scan*=0; this often results in track swapping. The value of the first tracking stage is that it removes significant number of extraneous contacts. Next, we break all contact associations, and provide the resulting cleaner set of contact data (with FAR close to zero) to a second tracking stage, now with *n-scan*>0.

As shown in figure 7, the results of the *track-break-track* approach are impressive. For all *n-scan* settings, we achieve a significant reduction in the critical angle for track swap, and at significantly lower computational expense. Interestingly, it appears that we continue to achieve further performance benefits with increasing *n-scan*, even as the single-stage architecture has reached saturation.
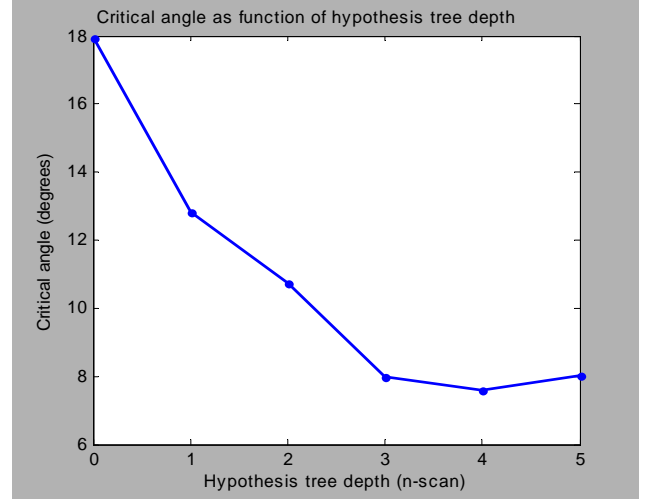

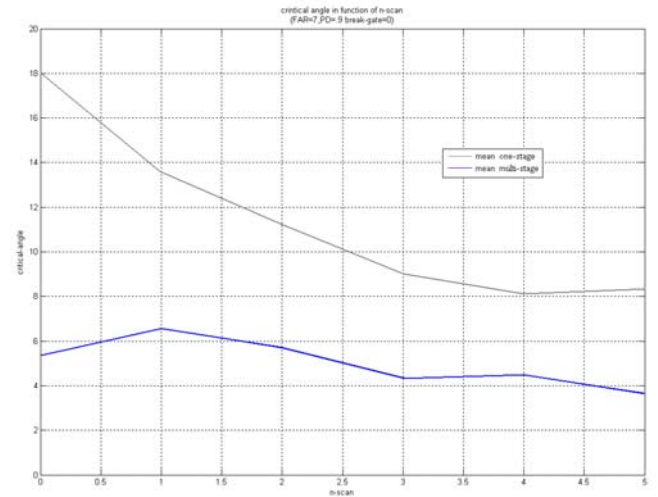Figure 6. Decreased track swapping with increasing hypothesis decision latency.

Figure 7. Lower critical angle for all n-scan in the track-break-track approach. Results are based on PD=0.9 and FAR=7.

6   Further Analysis

It is of interest to understand the track repulsion effect at a simple analytical level, particularly the sub-optimality of scan-based tracking with respect to optimal track estimation as achieved with a batch, *maximum likelihood* (ML) approach.

We consider a simple scenario. Two targets in one-dimensional Cartesian space are observed with linear measurements in Gaussian noise; for target $i \in \{1,2\}$, we

denote by $X_k^{(i)}$ the state at time k, and measurements are given by the following, where $v_k^{(i)} \sim N(0, \sigma^2)$ is a zero mean Gaussian random variable that is uncorrelated with other measurement errors:

$$Y_k^{(i)} = X_k^{(i)} + v_k^{(i)}, \ k = 1, ..., N . \qquad (1)$$

Assume that both targets are known to be stationary and at unknown locations equidistant from the coordinate system origin:

$$X_k^{(1)} = -X_k^{(2)} = \overline{X}, k = 1, ..., N . \qquad (2)$$

For simplicity, we neglect false contacts and missed detections. Thus, the measurement-origin uncertainty is limited to confusion as to which target gives rise to which measurement. Assume that a large number N of data scans is available.

It can be shown that, for $N \rightarrow \infty$, the ML solution is given by the following location estimates $\hat{X}^{(1)}$ and $\hat{X}^{(2)}$:

$$\hat{X}^{(1)} = -\hat{X}^{(2)} = \overline{X} . \qquad (3)$$

That is, the ML solution does not suffer any track-repulsion bias. On the other hand, for $N \rightarrow \infty$, the scan-based solution as obtained e.g. with a multi-hypothesis tracker with a constant position kinematic motion model is the following, where $v, w \sim N(0, \sigma^2)$ are uncorrelated random variables:

$$\hat{X}^{(1)} = -\hat{X}^{(2)} = E\big[\max\{\overline{X} + v, -\overline{X} + w\}\big]. \qquad (4)$$

The track repulsion effect is given by the displacement $\hat{X}^{(1)} - \overline{X}$, or $-\left(\hat{X}^{(2)} - \overline{X}\right)$. Figure 8 illustrates its magnitude for a range of values for target spacing ($2\overline{X}$), and assuming $\sigma = 10$ [m]. We see that the scan-based solution suffers a non-trivial bias for sufficiently close targets. For example, at a target spacing equal to the measurement error standard deviation of 10m, the track displacement is approximately 2m, i.e. the tracks are approximately 14m apart. This example helps to understand the same fundamental phenomenon leading to the track-swapping results in the previous section.
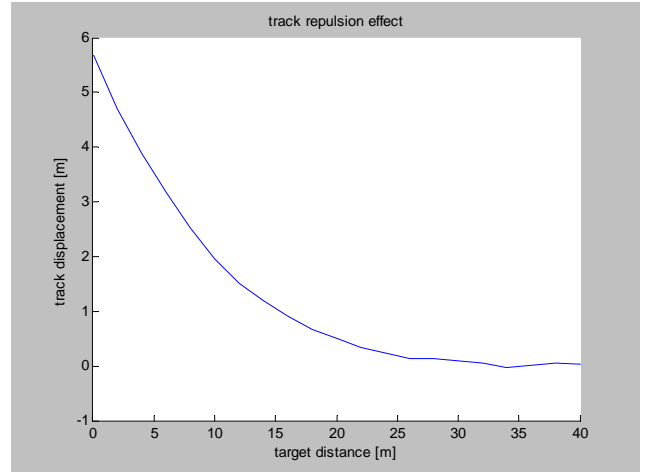

Figure 8. Track displacement as a function of target separation, for scan-based tracking of stationary targets.

7 Conclusions and Future Directions

We have found that the NURC DMHT in a *track-break-track* configuration is effective at combating the track repulsion effect in difficult multi-target scenarios. In particular, computation times are limited, allowing for large settings for the hypothesis tree depth. Tracking results outperform the single-stage approach, as demonstrated by the smaller critical angles for track swapping. Future work on this topic should include an investigation of a wider range of scenario settings, and a comparison with other tracking approaches including those documented in [10].

Additionally, the *track-break-track* approach should be explored in multi-sensor settings where complementary data is available from sensors with widely-varying update rates. Figure 9 illustrates this setting.
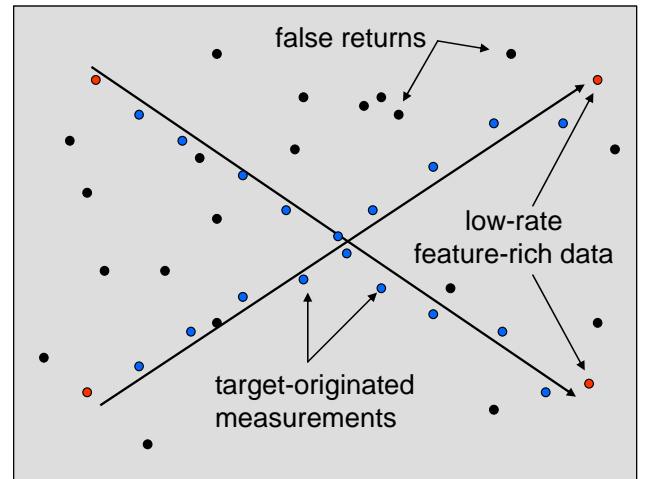

Figure 9. Multi-scale data for which the *track-break-track* approach holds promise.

As shown in the figure, assume that we have a crossing-target scenario with a primary sensor that provides rapid data scans. As seen earlier, it is difficult to determine

whether a target crossing has occurred, particularly if the target trajectories cross slowly relative to the scan rate. Assume that a secondary sensor provides highly informative target feature data, though with a low scan rate.

In such a scenario, use of the *track-break-track* approach should be explored. In particular, in the first tracking stage we would only process data from the primary sensor, with the objective of significantly reducing the false returns. In this first stage, we would have little confidence that successful tracking through the crossing will have been achieved. Subsequently, track labels are removed and the data is tracked again, this time with the inclusion of the feature-rich secondary sensor data, and with a large *n-scan* setting. The feature data coupled with the large *n-scan* setting will allow for successful determination of whether the targets have crossed or not.

Note that our approach as outlined here obviates the need for determining when association hypotheses are high-likelihood or not, which would be required in a cumbersome alternative two-stage approach to the problem whereby tracklet would be formed before and after the target crossing, followed by tracklet fusion with the inclusion of secondary sensor data. Note that the first-stage tracklet formation would utilize small association gates, i.e. there would necessarily be high fragmentation resulting from the first stage of tracking.

Finally, it is not clear whether a single-stage processing approach could effectively address this tracking problem. The required *n-scan* setting for effective association of secondary sensor returns would lead to computational intractability.

8   Acknowledgement

References

[1]   S. Blackman and R. Popoli, *Design and Analysis of Modern Tracking Systems*, Artech House, 1999.

[2]   D. Avitzour, A Maximum Likelihood Approach to Data Association, *IEEE Transactions on Aerospace and Electronic Systems*, vol. 28(2), April 1992.

[3]   D. Reid, An Algorithm for Tracking Multiple Targets, *IEEE Transactions on Aerospace and Electronic Systems*, vol. 24(4), December 1979.

[4]   T. Kurien, Issues in the Design of Practical Multi-Target Tracking Algorithms, in *Multi-Target Multi-Sensor Tracking: Advanced Applications*, Artech House, 1990.

[5]   S. Coraluppi and C. Carthel, Distributed Tracking in Multistatic Sonar, *IEEE Transactions on Aerospace and Electronic Systems*, vol. 41(3), July 2005.

[6]   C. Carthel, S. Coraluppi, and A. Maguer, Modeling and Validation of Maritime Surveillance Performance, in *Proceedings of the 3rd Annual Maritime Systems and Technology Conference*, November 2008, Cadiz, Spain.

[7]   S. Coraluppi and C. Carthel, Multi-Stage Data Fusion and the MSTWG TNO Datasets, submitted to the *12th International Conference on Information Fusion*, July 2009, Seattle WA, USA.

[8]   R. Fitzgerald, Track Biases and Coalescence with Probabilistic Data Association, *IEEE Transactions on Aerospace and Electronic Systems*, vol. 21(6), November 1985.

[9]   H. Blom and A. Bloem, Probabilistic Data Association Avoiding Track Coalescence, *IEEE Transactions on Automatic Control*, vol. 45(2), February 2000.

[10]   P. Willett, T. Luginbuhl, and E. Giannopoulos, MHT Tracking for Crossing Sonar targets, in *Proceedings of the SPIE Conference on Signal and Data Processing of Small Targets*, August 2007.

[11]   O. Gerard, C. Carthel, S. Coraluppi, and P. Willett, Feature-Aided Tracking for Marine Mammal Detection and Classification, *Canadian Acoustics*, vol. 36(1), March 2008.

[12]   B. Anderson and J. Moore, *Optimal Filtering*, Prentice-Hall, 1979.

[13]   D. Bertsekas, *Nonlinear Programming*, Athena Scientific, 1995.

[14]   S. Coraluppi, D. Grimmett, and P. de Theije, Benchmark Evaluation of Multistatic Sonar Trackers, in *Proceedings of the 9th International Conference on Information Fusion*, July 2006, Florence, Italy.